REJECTIONS UNDER 35 U.S.C. §102

Claims 1-31, 33-35, 37-49, and 41-42 are rejected under §102 as purportedly being anticipated by Pandya. These rejections are respectfully traversed.

I. Overview Of Embodiments Of The Invention

Embodiments of the invention are directed to a content addressable storage (CAS) system, which is one by which a unit of data stored on the CAS system is accessed using an address derived at least partially from the content of the unit of data. (specification, page 1, lines 10-12). The CAS system can serve as storage for a host computer. When a host computer sends a request to the CAS system to retrieve a unit of data, the host provides the content address of the unit of data, and the storage system then determines, based on the content address the physical location where the unit of data is stored so that it can be retrieved and returned to the host computer.

The task of determining the physical location for a unit of data may have several aspects, particularly when the storage system is a distributed storage system made up of a number of separate nodes. (page 1, lines 22-23). To determine the physical location of a unit of data on such a storage system, the storage system first determines on which node the unit of data is stored, and then determines which disk on that node the data is stored. (page 1, lines 26-28).

One known method of determining which storage nodes stores a particular unit of data is referred to as a multi cast location query (MLQ) (page 2, lines 9-10). Using this technique, a message is broadcast to each storage node that stores units of data, asking if it stores the particular unit of data to be accessed. (page 2, lines 10-13). Each storage node then determines if it stores the requested unit of data, and may do so by accessing a database or table that lists the units of data stored thereby. (page 2, lines 13-17). An MLQ is a computationally expensive process, as it requires each storage node to perform an exhaustive database search for each unit of data requested. (page 3, lines 15-16).

To reduce the computational expense of using an MLQ to locate units of data on the storage system, one embodiment of the invention is directed to a technique (referred to in the specification as intrinsic locations) that employs at least a portion of the content address to

determine at least one aspect of the storage location for units of data on the storage system. (page 24, lines 27-30). In one example of Fig. 5, various storage nodes are assigned content addresses beginning with particular characters. (page 25, lines 1-6). By selecting the physical storage location of the unit of data based on the content address, the storage system may determine the physical location later (e.g., when a request is received to read the content unit) without a database look up or a multi-cast location query, as the storage system can simply examine the content address and determine the location of the unit of data based on the content address.

The foregoing overview is provided merely to assist the Examiner in appreciating various aspects of the present invention. The overview may not apply to each of the independent claims, and the language of the independent claims may differ in material respects from the overview provided above. The Examiner is requested to give a careful consideration to the language of each of the independent claims and to address each on its own merits, without relying on the overview provided above. Applicants do not rely on the overview to distinguish any of the claims of the present invention over the prior art, but rather, rely only upon the arguments provided below.

II. Discussion of Pandya

Pandya is directed to a high performance hardware IP processor that is used to relieve the performance impact on a host processor of implementing the TCP/IP stack in software (Pandya, ¶0067). Figure 17 of Pandya is an architectural diagram of the IP processor (Pandya, ¶0107). Pandya discloses, in connection with Figure 17, that input queue block 1701 queues packets as they arrive and operates in conjunction with packet scheduler 1702, which retrieves packets from the input queue and passes them for classification to classification engine 1703 (Pandya, ¶0108). Classification engine 1703 receives a packet, classifies the packet, and provides a classification tag to the packet before it is provided to processor array 1706(a)...1706(n) (Pandya, ¶0108).

Figure 20 is more detailed diagram of classification engine 1703. Pandya discloses that the classification engine examines various fields of the received packet to identify the type of packet, the protocol type (e.g., IP, ICMP, TCP, UDP, etc.), the port addresses, the source and destination fields, and other fields (Pandya, ¶0118). Pandya discloses that the classifier, in some

embodiments, may use a content addressable memory (CAM) array to assist in identifying packets (Pandya, ¶0118). As Applicants explained during the telephone interview, the classifier passes packet fields to the CAM to determine if the packet fields match any of the values stored in the CAM array(Pandya, ¶0118). When there is a match, the CAM provides an action/even tag. All of the action/event tags for a packet are compiled and are used to generate a classification tag for the packet, which may indicate the flow or session ID, the protocol type, or other classification information (Pandya, ¶0118).

III. Claims 1-10

Claim 1 is directed to a method of processing data in a computer system comprising at least one host and at least one content addressable storage (CAS) system, wherein the at least one host accesses data units stored on the at least one storage system using content addresses generated based on the content of the data units. The method comprises an act of, in response to an access request from the at least one host for a unit of data identified by a content address, parsing the content address to determine at least one aspect of a physical storage location for the unit of data on the at least one storage system.

Pandya does not disclose or suggest a computer system comprising a content addressable storage system that stores data for at least one host computer. The Office Action states that Pandya discloses a content addressable storage system in the form of the content addressable memory, disclosed in paragraph 0118, that forms a part of classification engine 1703 (*see* Office Action, page 2, ¶2). Applicants disagree that the content addressable memory is a "content addressable storage system which stores data for the at least one host," as required by claim 1.

As discussed during the telephone interview, the content addressable memory of Pandya does not store data for a host computer. The content addressable memory of Pandya is a memory that is preprogrammed with packet fields and values (Pandya, ¶0118, lines 35-40). When network packets are received, data in the packets are compared to the values that are pre-stored in the content addressable memory so that the packet may be classified. **Pandya does not disclose or suggest that the content addressable memory stores packets provided by a host**

computer or retrieves packets from the content addressable memory and provides them to the host computer.

Further, Pandya does not disclose or suggest that data units stored on the storage system (allegedly the CAM) are accessed using a content address generated based on the content of the data units. The Office Action asserts that Pandya discloses this feature at ¶0121, lines 76-100. However, the cited portion of Pandya does not even mention a content address, let alone a host accessing a data unit using the content address.

To be consistent with the rest of the rejection, the claimed content addressable storage system must purportedly read on the CAM array. However, as Applicants explained during the telephone interview, no matter what entity in Pandya is considered to be the host, Pandya does not disclose or suggest any entity accessing a data unit in the CAM array using a content address. The CAM array of Pandya serves as a lookup table which is used to classify a packet received from the host. The "address" used to access the CAM array is derived from the bits of the packet (Pandya, ¶0118). To the extent that it can be asserted that any "host computer" is accessing classification information stored in the CAM array, this information is not accessed using a content address computed from the content of the classification information stored in the content addressable memory. Pandya simply does not disclose or suggest a host computer accessing a data unit on a content addressable storage using a content address **that is computed from the content of the data unit**.

Additionally, Pandya does not disclose or suggest that, in response to an access request for a content unit, a content address is parsed to determine at least one aspect of a physical storage location for the unit of data on the storage system. The Office Action asserts that Pandya discloses this limitation at ¶0124, lines 3-8, which states, "[t]he storage flow and RDMA controller block provides the functionality necessary for the host to queue the commands (storage or RDMA or sockets direct or a combination thereof) to this processor, which then takes these commands and executes them, interrupting the host processor primarily on the command termination." It is unclear how the Office Action interprets this paragraph to disclose parsing a content address that is computed from the content of a data unit to determine at least one aspect of a physical storage location for the data unit, as this paragraph mentions neither using a content

address for a data unit, parsing the content address, nor determining a storage location for the data unit. As should be clear from the discussion above, Pandya does not disclose or suggest the use of a content address for a data unit that is computed, at least in part, from the content of the data unit. It follows that Pandya does not disclose or suggest determining at least one aspect of physical storage location for a unit of data by parsing such a content address.

In view of the foregoing, claim 1 patentably distinguishes over Pandya. Accordingly, it is respectfully requested that the rejection of claim 1 under 35 U.S.C. §102(e) be withdrawn.

Claims 2-10 depend from claim 1 and are patentable for at least the same reasons. Accordingly, it is respectfully requested that the rejection of these claims be withdrawn.

IV. Claims 11-20

Claim 11 is directed to at least one computer readable medium encoded with instructions that perform a method substantially similar to that recited in claim 1.

Thus, claim 11 patentably distinguishes over Pandya for the same reasons discussed above in connection with claim 1. Accordingly, it is respectfully requested that the rejection of claim 11 under 35 U.S.C. §102(e) be withdrawn.

Claims 12-20 depend from claim 11 and are patentable for at least the same reasons. Accordingly, it is respectfully requested that the rejection of these claims be withdrawn.

V. Claims 21-30

Claim 21 is directed to a content addressable storage system for use in a computer system, including the content addressable storage system and at least one host, wherein the at least one host accesses data units stored on the content addressable storage system using content addresses generated based on the content of the data units. The content addressable storage system comprises:

at least one storage device to store data received from the at least one host; and at least one controller that, in response to an access request from the at least one host computer for a unit of data identified by a content address, parses the content address to determine at least one aspect of a physical storage location for the unit of data on the at least one storage system.

As should be clear from the discussion above, Pandya does not disclose or suggest that at least one host computer accesses data units stored on the content addressable storage system using a content address generated based on the content of the data units, and does not disclose or suggest that, in response to an access request for a content unit, the content addressable storage system parses the content address to determine at least one aspect of a physical storage location for the unit of data on the storage system.

Thus, claim 21 patentably distinguishes over Pandya. Accordingly, it is respectfully requested that the rejection of claim 21 under 35 U.S.C. §102(e) be withdrawn.

Claims 22-30 depend from claim 21 and are patentable for at least the same reasons. Accordingly, it is respectfully requested that the rejection of these claims be withdrawn.

VI. Claims 31-34

Claim 31 is directed to a method of processing data in a computer system comprising at least one host and at least one content addressable storage system which stores data for the at least one host, wherein the at least one host accesses data units stored on the at least one storage system using content addresses generated based on the content of the data units. The method comprising acts of: (a) receiving, from the host, a request to store a unit of data on the storage system, the unit of data having a content address based on the content of the unit of data; (b) determining, based on the content address, a first storage location on the storage system to which the content address maps;
(c) storing a pointer for the first unit of data at the first storage location, the pointer pointing to a second storage location; and (d) storing the unit of data at the second storage location on the storage system.

As should be clear from the discussion above, Pandya does not disclose or suggest a computer system comprising a content addressable storage system that stores data for at least one host computer, and does not disclose or suggest a content addressable storage system that receives, from the host, a request to store a unit of data having a content address based on the content of the unit of data.

In addition, as Pandya does not disclose or suggest the use of content addresses to identify units of data, Pandya fails to disclose determining, based on a content address, a first storage location on the storage system to which the content address maps.

Further, Pandya fails to disclose or suggest, "storing a pointer for the first unit of data at the first storage location, the pointer pointing to a second storage location" and "storing the unit of data at the second storage location on the storage system," as recited in claim 31. The Office Action asserts that Pandya discloses storing a pointer for the first unit of data at the first storage location at ¶0121, lines 27-29. However, this portion of Pandya does not even mention storing a pointer, let alone a pointer that is stored at a storage location determined from the content address of the data unit.

As Pandya fails to disclose any of the limitations discussed above, claim 31 patentably distinguishes over Pandya. Accordingly, it is respectfully requested that the rejection of claim 31 under 35 U.S.C. §102(e) be withdrawn.

Claims 32-34 depend from claim 31 and are patentable for at least the same reasons. Accordingly, it is respectfully requested that the rejection of these claims be withdrawn.

VII. Claims 35-38

Claim 35 is directed to at least one computer readable medium encoded with instructions that, when executed on a computer system, perform a method substantially similar to that recited in claim 31.

As should be clear from the discussion above, Pandya fails to disclose any of the limitations of claim 35. Thus, claim 35 patentably distinguishes over Pandya. Accordingly, it is respectfully requested that the rejection of claim 35 under 35 U.S.C. §102(e) be withdrawn.

Claims 36-38 depend from claim 31 and are patentable for at least the same reasons. Accordingly, it is respectfully requested that the rejection of these claims be withdrawn.

VIII. Claims 39-42

Claim 39 is directed to a content addressable storage system for use in a computer system that includes at least one host, wherein the at least one host accesses data units stored on the

content addressable storage system using content addresses generated based on the content of the data units. The content addressable storage system comprises: at least one storage device to store data received from the at least one host; and at least one controller that: receives, from the host, a request to store a unit of data on the storage system, the unit of data having a content address based on the content of the unit of data; determines, based on the content address, a first storage location on the storage system to which the content address maps; stores a pointer for the first unit of data at the first storage location, the pointer pointing to a second storage location; and stores the unit of data at the second storage location on the storage system.

As should be clear from the discussion above, Pandya fails to disclose any of the limitations of claim 39. Thus, claim 39 patentably distinguishes over Pandya. Accordingly, it is respectfully requested that the rejection of claim 39 under 35 U.S.C. §102(e) be withdrawn.

Claims 40-42 depend from claim 39- and are patentable for at least the same reasons. Accordingly, it is respectfully requested that the rejection of these claims be withdrawn.

REJECTIONS UNDER 35 U.S.C. §103

Several dependent claims (i.e., claim 32, 36, and 40) are rejected under 35 U.S.C. §103(a) as purportedly being obvious over Pandya. These rejections are respectfully traversed. Applicants do not address these rejections in detail, because each of the independent claims from which these claims depend is patentable over Pandya for the reasons discussed above.

The lack of comment on the rejection of the dependent claims under 35 U.S.C. §103(a) should not be construed as acquiescence to the analysis applied in rejection these claims. For example, Applicants disagree with the assertion in the Office Action that certain of the claims that recite an order of actions are rendered obvious even though the prior art relied upon is not alleged to show such an order. *See* Office Action, page 20.

CONCLUSION

In view of the foregoing amendments and remarks, this application should now be in condition for allowance. A notice to this effect is respectfully requested. If the Examiner believes, after this amendment, that the application is not in condition for allowance, the Examiner is requested to call the Applicant's attorney at the telephone number listed below.

If this response is not considered timely filed and if a request for an extension of time is otherwise absent, Applicant hereby requests any necessary extension of time. If there is a fee occasioned by this response, including an extension fee, that is not covered by an enclosed check, please charge any deficiency to Deposit Account No. 23/2825.

Dated: April 11, 2006

Respectfully submitted,

Michael Kilian et al., Applicants

By _____

Richard F. Giunta
Registration No.: 36,149
WOLF, GREENFIELD & SACKS, P.C.
Federal Reserve Plaza
600 Atlantic Avenue
Boston, Massachusetts 02210-2206
(617) 646-8000